

ABSTRACT

The present invention relates to a system and method of optimizing execution of a distinct count query. The system and method allows clients or database administrators to improve queries by properly designing data cubes and partitions of the data in the cube.

5 The partition data can also be ordered so as to facilitate determining the range of each partition. Partitions with overlapping ranges can be executed in parallel. Furthermore, partitions with non-overlapping ranges can also be executed in parallel to optimize query execution rather than digressing from parallel to sequential execution by virtue of their range.

10